arXiv:1507.02973v1 [cs.SI] 10 Jul 2015

Overcoming Data Scarcity of Twitter: Using Tweets as Bootstrap with Application to Autism-Related Topic Content Analysis

Adham Beykikhoshk† Ognjen Arandjelović‡ Dinh Phung Svetha Venkatesh
Centre for Pattern Recognition and Data Analytics
School of Information Technology
Deakin University
Geelong, VIC 3216
Australia
{†adham.beyki,‡ognjen.arandjelovic}@gmail.com

***Abstract*—Notwithstanding recent work which has demonstrated the potential of using Twitter messages for content-specific data mining and analysis, the depth of such analysis is inherently limited by the scarcity of data imposed by the 140 character tweet limit. In this paper we describe a novel approach for targeted knowledge exploration which uses tweet content analysis as a preliminary step. This step is used to bootstrap more sophisticated data collection from directly related but much richer content sources. In particular we demonstrate that valuable information can be collected by following URLs included in tweets. We automatically extract content from the corresponding web pages and treating each web page as a document linked to the original tweet show how a temporal topic model based on a hierarchical Dirichlet process can be used to track the evolution of a complex topic structure of a Twitter community. Using autism-related tweets we demonstrate that our method is capable of capturing a much more meaningful picture of information exchange than user-chosen hashtags.**

I. Introduction

The last decade has witnessed the emergence of a number of social media platforms and the rapid rate of their proliferation, breadth of adoption, and frequency of use. The associated wealth of user-generated content, most of which is easily accessible, has created an immense opportunity for automatic data mining and knowledge discovery. It is unsurprising that this opportunity has been readily embraced both by the academic community [8], [12], [19], [25], [27], [36], [42] and the commercial sector [29], [37], [44]. Nevertheless research in this area is still in its infancy [7], with the majority of the methodologies applied to the processing of social media content to date being simple adaptations of existing algorithms developed for the analysis of 'conventional' data sources. At this stage neither the full range of the associated challenges nor the potential of information found on social media are fully understood, and only when these tasks are accomplished will the development of algorithms specifically aimed at social media be able to achieve substantial progress.

The backdrop we just summarized readily explains why Twitter in particular has attracted so much research attention, as illustrated by the plot in Fig. 1. Firstly, it is immensely popular, with 288 million active users and 500 million tweets sent per day, with a high level of global adoption witnessed by the statistic that 77% of its users are located outside the USA [51]. Secondly, given that research in this field is still in its early stages, the purely textual nature of content on Twitter itself provides a sufficiently constrained environment needed to develop an understanding of the available information and the performance of different data mining and machine learning methodologies. In large part these observations motivate the focus of our work on Twitter too.

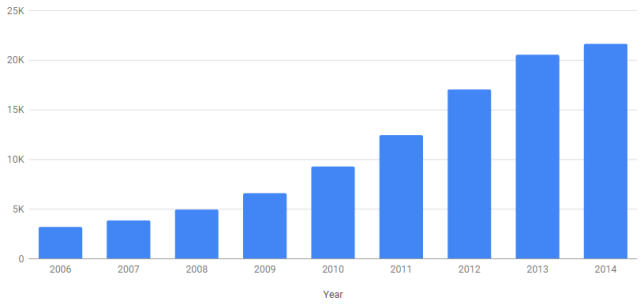

Fig. 1. Number of articles indexed by Google Scholar matching the search query *twitter mining*, shown per year of publication (note that the count is *per annum*, not cumulative).

The specific interest of the present work lies in the use of Twitter content for the monitoring and analysis of attitudes, behaviours, and conversations pertaining to health care and the autism spectrum disorder (ASD) [1] in particular. Our interest in this condition is motivated by two key factors. The first of these concerns the associated human and economic cost [32]. In brief, ASD is a life-long neurodevelopmental disorder [35] characterized by severe impairments in social interaction, communication, and in some cases cognitive abilities [38]. ASD typically begins already in infancy or at the very latest by the age of three when it is usually detected by an abnormal

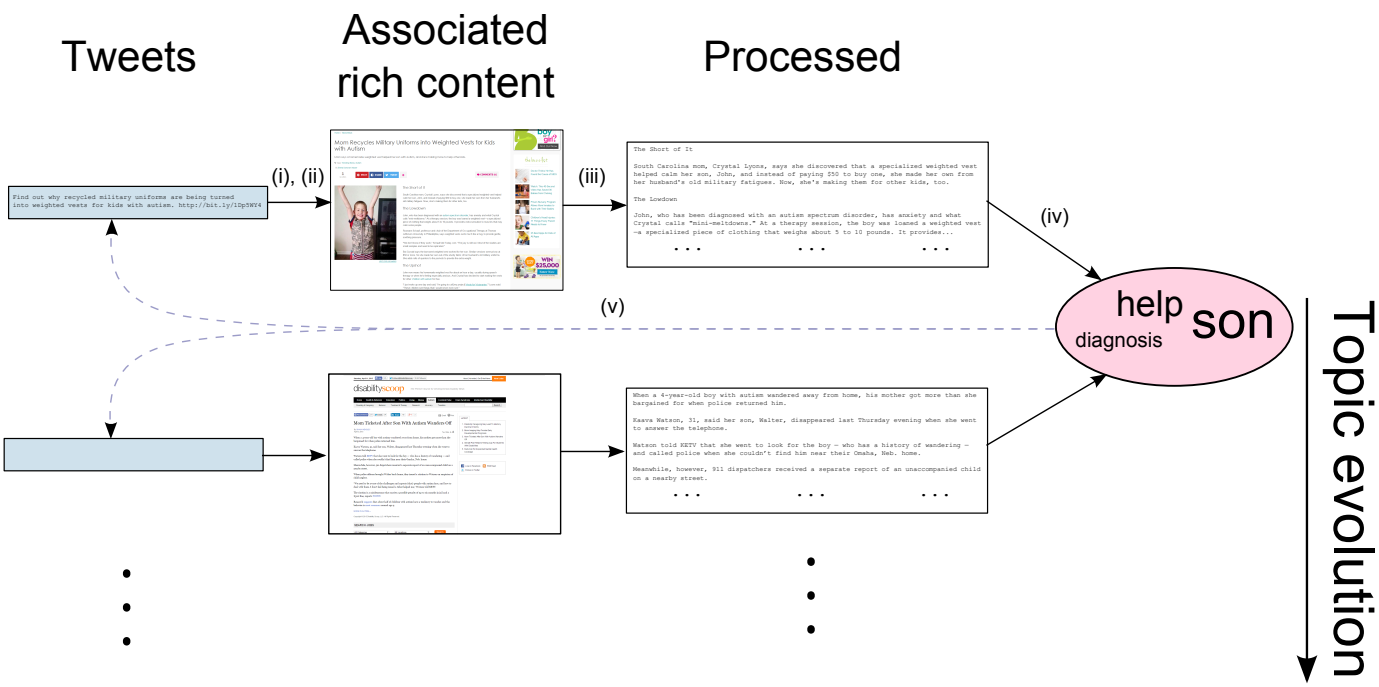

Fig. 2. Conceptual overview of the proposed framework. Our analytical pipeline consists of: (i) identifying URLs in tweets, (ii) following the corresponding links, (iii) processing the target web pages to extract relevant textual information, (iv) performing topic analysis on these documents, and finally (v) linking the extracted topic back with the original tweets.

lack of social reciprocity. Current evidence suggests that a stable proportion of approximately 0.5–0.6% of the population is afflicted by ASD although the actual diagnosis rate is on the increase due to the broadening diagnostic criteria [9].

The second important factor motivating our research stems from the potential of Twitter as a pervasive interactive platform to reach ASD communities which are notoriously difficult to engage by traditional means. This includes both those segments of the ASD community which are isolated socially, for example due to stigma often attached to mental disorders [22], as well as those which are isolated geographically, for example in remote rural communities. From the public policy point of view, understanding the practices and beliefs of parents and carers of ASD-affected individuals is crucial, yet often lacking [24]. The penetration of beliefs, and behavioural and educational interventions which are questionable [53] and poorly supported by evidence (e.g. gluten-free and casein-free diets, and cognitive behavioural therapy [20]), and sometimes outright in conflict with science [50] is particularly worrying [53].

Thus the key premise behind our research efforts is that the rapid rise in the adoption of social media as a platform for the expression and exchange of ideas, which facilitates the emergence of special interest communities, can be used to study and monitor the beliefs and practices of the population affected by ASD. Considering the challenge of reaching and engaging with this specific target population, success in this domain of data mining would be a powerful utilization of social networks for public good.

A. *Related previous work*

Most of the previous work on data mining Twitter has not been automatic, that is, analysis was performed 'manually' by humans within the proof-of-concept paradigm. Considering the usage statistics summarized in Sec. I, it is evident that non-machine-based approaches are not feasible in practice – the process is extremely laborious which severely limits how much data can be processed. A major challenge in automating the analysis lies in the likely semantic gap needed to interpret Twitter messages. As we shall demonstrate in the next section, the severe limit on the character length of tweets appears to exaggerate this gap yet further, requiring an even more sophisticated understanding of language, its idioms, sociolects etc. The following ASD-related tweets, representative of the corpus we collected (see Sec. III), readily illustrate this:

Example tweet 1:

Looks like we will have more #autism research happening for children in #EarlyIntervention next year! :-) #VisualSupports #MobileTech

Example tweet 2:

Authors who see autism as "tremendously burdening" elicit dire views of autism from parents http://j.mp/1FyqXwF "Ethical approval: none"

Example tweet 3:

101 autism: Genetic analysis of individuals with autism finds gene deletions - Using powerful genetic sequencing. http://is.gd/UhprQK

Example tweet 4:

#Apple #Censorship & Dr. Brian Hooker Interview exposing CDC Cover-up of the Vaccine & Autism Link on .@rediceradio http://youtu.be/19uvPtg6SPI

Although of limited practical use in the long term, manual approaches to Twitter data analysis and mining have served an important role in illustrating the potential of this platform. Indeed the range of applications in which the usefulness of Twitter data has been demonstrated is remarkable. For example, the inference of message sentiment from emoticons was described by Jiang *et al.* [34], and Bifet and Frank [13], while Agarwal *et al.* [3] employed linguistics-based classifiers to a similar end. Sentiment analysis in tweets was further used by Bollen *et al.* [15] in the prediction of socio-political, cultural, and economic events. Sociometric applications were also explored by Mitchell *et al.* [39]. The potential of Twitter in the understanding and management of emergency situations, such as earthquakes, attracted significant research attention too [52], with successes reported by Robinson *et al.* [43] and Sakaki *et al.* [45] amongst others.

The same themes drawing from the use of geographical information and the real-time nature of Twitter can be observed in the existing work on the use of the platform in the realm of health care [17]. For example Paul and Dredze [41] tracked over time the spread of illnesses, changes in behavioural risk factors, and symptoms and patterns of medication use. Patterns of antibiotic use were analysed by Scanfeld *et al.* [47] while Jashnisky *et al.* [33] studied the possibility of predicting suicide rates.

Notwithstanding the increasingly recognized need for more effective information exchange with the affected communities, the use of social media in the understanding and management of ASD has received little research attention. A manual (i.e. human effort-based, rather than computerized) analysis of ASD-affected individuals' writing patterns was reported by Newton *et al.* [40] while our previous work [11] described the only automatic system in this domain.

II. Proposed method

In this section we lay out the key contributions of the present work. We start with an overview of the proposed methodology and then proceed with a detailed description of each of its components in turn.

A. Data mining framework – key ideas

In this paper our aim is to extract the topic content of a specific category of tweets and track its changes over time. Although our focus is on ASD-related tweets, which we showed can be readily detected automatically [11], the nature of the proposed framework is entirely general and could in principle be applied to any other category of tweets. Here 'topic' (and any derivations thereof) is used as a technical concept from the field of text analysis, that is, it refers to a probability distribution over a collection of words/terms (the 'dictionary') [14]. Intuitively, this can be interpreted as a formalization of the colloquial meaning of the word 'topic' [16].

Probably the most obvious approach to tweet topic modelling involves the use of Twitter's built-in capability for tagging tweets by their authors, using so-called 'hashtags'. These can be recognized by the leading special character '#' (the hash sign) and are understood to be meta-data labels. We found that approximately 25% of ASD-related tweets contain hashtags. However recent research has shown that a small number of keywords is inadequate in capturing the essence of the content of the associated tweet [14].

The idea that the main textual information within tweets can be analysed for its topic content is motivated by the findings reported in our previous work [11]. In particular we showed that even after terms which are *ipso facto* ASD-related are removed, there remains sufficient information (both structural and semantic) which is characteristic to ASD-related tweets. However our attempts at applying a variety of different topic models failed to produce meaningful results. Specifically we found that the number of terms contained in individual tweets (the average term count per tweet in our corpus is 13) is too small to give rise to topics of sufficient complexity. Considering that this bottleneck is inherent to Twitter rather than a limitation of our methodology, at first sight it appears that topic content extraction from tweets is impossible, to say nothing of any temporal analysis thereof. However, in this work we show that this barrier can be overcome through an ingenious trick. In particular we observed that a large number of tweets contain URLs i.e. links to external content, as illustrated in Fig. 3. Moreover these tweets are exactly those that we are most interested in, that is, those which aim to contribute the greatest amount of information.

Hence our idea is to use tweets not as endpoint data mining sources themselves but rather as data mining intermediaries used to discover much richer associated content. Some previous work has already recognized the potential of such external content e.g. [46], [2]. However, unlike Sankaranarayanan *et al.* [46] who use linked content for clustering and the identification of popular news topics, or Abel *et al.* [2] who focus on building news-related user models, we demonstrate that content harvested from linked URLs can be used to build a complex, temporal and cross-modal graph which facilitates temporal topic analysis which can be mapped back to Twitter posts (an interesting comparison can be made with machine learning methods in other fields which exploited such parallel structures with success e.g. for inferring pose-wise mappings across face manifolds [6], [4]). In brief, our analytical pipeline consists of: (i) identifying URLs in tweets, (ii) following the corresponding links, (iii) processing the target web pages to extract relevant textual information, (iv) performing topic analysis on these documents, and finally (v) linking the extracted topic back with the original tweets. This is illustrated conceptually in Fig. 2. Our automatic extraction of relevant textual information from web pages is described in the next section. A method for tracking complex temporal changes of the topic structure of a longitudinal document

corpus, based on a graph constructed over topics extracted using a hierarchical Dirichlet process-based model for different epochs is explained in Sec. II-C.

B. Automatic data enrichment using URLs

As indicated in the previous section, the key idea of the proposed method for Twitter topic analysis and tracking is to map tweets which contain URLs to the space of the corresponding web pages. This process effectively establishes correspondence between a tweet and the web page linked through a URL in the tweet body. Observe that in principle a single tweet can be put in correspondence with multiple web pages as multiple URLs may be contained within the tweet, though in practice this is seldom the case – we found that over 96% of URL-containing tweets include a single URL, and fewer than 0.1% more than two URLs. Equally, a single web page can be linked to from different tweets; this is a far more frequent occurrence. Whatever the exact network of correspondences may be, topics assigned to the content of a specific web page are also associated with all tweets which contain a URL link to this web page. As per our previous observations, this means that a tweet with multiple URLs 'inherits' topics from multiple web pages, and equally, a web page lends its topics to all tweets which link to it.

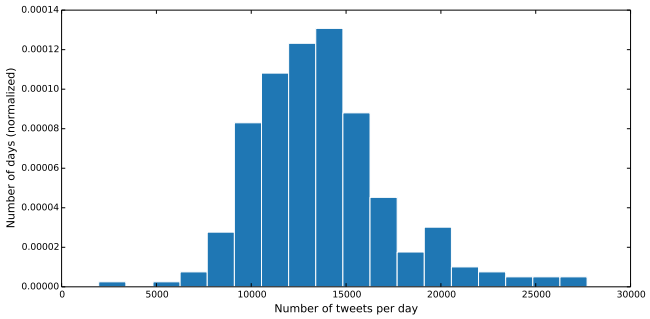

(a) Daily tweeting rate

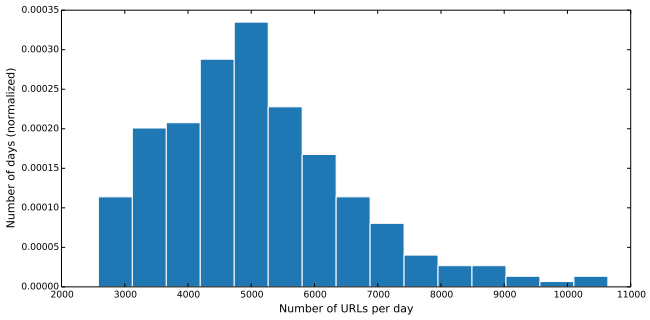

(b) Daily in-tweet URL rate

Fig. 3. The statistics of relevant data availability rate, estimated from our data set (please see Sec. III for additional detail). Shown are histograms corresponding to (a) the daily number of tweets identified as ASD-related [11], and (b) the daily count of URLs harvested from these tweets.

We retrieve tweets using Twitter's stream and search APIs (please see Sec. III for further detail). This allows us to fetch relevant tweet information, such as the text of the message, posting time and date, geo-tags, and the author's user ID, as well as the list of any URLs linked to from the tweet in the JavaScript Object Notation (JSON) format. Therefore no sophisticated processing of the tweet is needed at this stage – the URLs of interest are retrieved explicitly. These are then used to fetch the corresponding web pages. We process only those URLs which link to HTML pages. For example by examining the 'Content-Type' response from the HTTP server we immediately reject links to videos and other non-text-based content. Pages which satisfy our validity criteria are then processed to extract their main text content. This is achieved by analysing the corresponding HTML descriptions and removing chunks of it which by their HTML tags can be identified as not being of interest. For example we remove text enclosed within tags such as `<COMMENT>`, `<META>`, `<HEADER>`, `<MENU>`, `<RSS>`, `<SPONSOR>`, and a series of others. Formatting tags such as `<FONT>`, `<DIV>`, `<P>`, and similar, are merely removed. We actively retain text within tags such as `<MAIN>`, `<ARTICLE>`, and `<BLOG>`, amongst others. This processes allows us to extract relevant content reliably without any manual input, with few errors. We found that rare errors do not affect the ultimate result of our method due to their randomness – there is no reason why any particular erroneous term (word) would consistently be included in the extracted text, which means that any *specific* error ends up being treated as noise in our bag-of-words model used as input to the main topic modelling and tracking algorithm described next.

C. Modelling topic evolution over time

Twitter is by its very nature highly dynamic: new content is rapidly created, transient periods of high popularity of specific ideas are followed by their rapid decline and the emergence of follow-up etc. Therefore it is imperative for the analysis of the topic structure of a tweet corpus to have a temporal component, which additionally has to be sufficiently robust and flexible to capture the range of complex interactions and causal cascades that the corresponding exchanges exhibit.

Our approach begins by discretizing time into epochs thereby dividing the tweet corpus into a series of chronological but overlapping sub-corpora. The topic structure of each sub-corpus associated with a particular epoch is then extracted separately using an HDP-based model. In other words each epoch is considered to span a sufficiently short time period that the associated sub-corpus of tweets can be treated as a static collection of documents. To ensure the validity of this assumption it is crucial that the time span of an epoch is chosen based on the speed of changes which the model is supposed to track (to reflect the highly dynamic nature of Twitter in this work we used three day epochs, with an overlap of successive epochs of two days). To associate topics across different epochs we create a similarity graph, connecting topics (as nodes) across successive epochs, and infer complex dynamics based on the strengths of these connections. This approach, first proposed in [10], is explained next.

1) Measuring inter-topic similarity: The key idea behind our approach stems from the observation that while topics may change significantly over time, by their very nature the change between successive epochs is limited. Therefore we infer the continuity of a topic in one epoch by relating it to all topics in the immediately subsequent epoch which are sufficiently similar to it under some similarity measure. This can be seen to lead naturally to a similarity graph representation whose nodes correspond to topics and whose edges link those topics in two epochs which are related. Formally, the weight of the directed edge that links $\phi_{j,t}$, the j-th topic in epoch t, and $\phi_{k,t+1}$ is set equal to $\rho\left(\phi_{j,t}, \phi_{k,t+1}\right)$ where ρ is an appropriate similarity measure. In this work we adopt the use of the Jaccard similarity between the probability distributions described by $\phi_{j,t}$ and $\phi_{k,t+1}$ i.e. $\rho\left(\phi_{j,t}, \phi_{k,t+1}\right) \equiv d_{\text{Jaccard}}\left(\phi_{j,t}, \phi_{k,t+1}\right)$.

A conceptual illustration of a similarity graph is shown in Fig. 4. It shows three consecutive time epochs $t-1, t$, and $t+1$ and a selection of topics in these epochs. Graph edge weight i.e. inter-topic similarity is encoded by varying the thickness of the corresponding line connecting two nodes – a thicker line signifies more similar topics. We use a threshold to eliminate weak edges automatically, retaining only the edges which correspond to sufficiently similar topics in adjacent epochs. It can be seen that this readily allows us to detect the disappearance of a particular topic, the emergence of new topics, as well as the splitting or merging of different topics. In summary:

Emergence If a node does not have any edges incident to it, the corresponding topic is taken as having emerged in the associated epoch (e.g. ϕ_{j+2} at time t in Fig. 4).

Disappearance If no edges originate from a node, the corresponding topic is taken to vanish in the associated epoch (e.g. ϕ_j at time t in Fig. 4).

Splitting If more than a single edge originates from a node, the corresponding topic is understood as being split into multiple topics in the next epoch (e.g. ϕ_i is split into ϕ_j and ϕ_{j+1} in Fig. 4).

Merging If more than a single edge is incident to a node, the topics of the nodes from which the edges originate are understood as having merged together to form a new topic (e.g. ϕ_i and ϕ_{i+1} merge to form ϕ_{j+1} in Fig. 4).

2) Hierarchical Dirichlet process mixture models: The Dirichlet process is a widely used prior for mixture modelling which readily allows a document collection to accommodate an arbitrary number of topics. It is the building block of many Bayesian nonparametric methods [21]. Succinctly put, a Dirichlet process $\mathrm{DP}\left(\gamma, H\right)$ is defined as a distribution of a random probability measure G over a measure space $(\Theta, \mathcal{B}, \mu)$, such that for any finite measurable partition $\left(A_1, A_2, \ldots, A_r\right)$ of Θ the random vector $\left[G\left(A_1\right), \ldots, G\left(A_r\right)\right]^T$ is a Dirichlet distribution with parameters $\left[\gamma H\left(A_1\right), \ldots, \gamma H\left(A_r\right)\right]^T$.

Owing to the discrete nature and infinite dimensionality

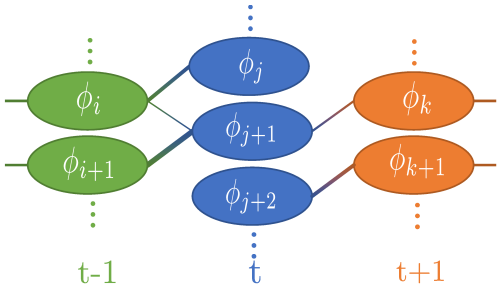

Fig. 4. A conceptual illustration of a similarity graph that models topic dynamics over time. A node corresponds to a topic in a specific epoch; the weight of an edge connecting two nodes is equal to the similarity of the corresponding topics.

of its draws, the DP is a highly useful prior for Bayesian mixture models. By associating different mixture components with atoms ϕ_k of the so-called stick-breaking process [31], and assuming $x_i|\phi_k \overset{iid}{\sim} F\left(x_i|\phi_k\right)$ where $F\left(.\right)$ is the likelihood kernel of the mixing components, the Dirichlet process mixture model (DPM) can be formulated. The DPM is suitable for nonparametric clustering of exchangeable data in a single group e.g. words in a document where the DPM models the underlying structure of the document with potentially an infinite number of topics. However, many real-world problems are more appropriately modelled as comprising multiple groups of exchangeable data (e.g. a collection of documents). In such cases it is usually desirable to model the observations of different groups jointly, allowing them to share their generative clusters. This idea is known as the sharing statistical strength and is achieved using a hierarchical structure.

Amongst different ways of linking group-level DPMs, HDP [48] offers an interesting solution whereby base measures of document-level DPs are drawn from another DP. In this way the atoms of the corpus-level DP (i.e. topics in our case) are shared across the corpus. Formally, if $\mathbf{x} = \{\mathbf{x}_1, \ldots, \mathbf{x}_J\}$ is a document collection where $\mathbf{x}_j = \left\{x_{j1}, \ldots, x_{jN_j}\right\}$ is the j-th document comprising N_j words, each document is modelled with a DPM $G_j|\alpha_0, G_0 \overset{iid}{\sim} \mathrm{DP}\left(\alpha_0, G_0\right)$ where its DP prior is further endowed by another DP $G_0|\gamma, H \sim \mathrm{DP}\left(\gamma, H\right)$. Since the base measure of G_j is drawn from G_0, it takes the same support as G_0. Also the parameters of the group-level mixture components, θ_{ji}, share their values with the corpus-level DP support on $\{\phi_1, \phi_2, \ldots\}$. The posterior for θ_{ji} follows a Chinese restaurant franchise process which can be used to develop inference algorithms based on Gibbs sampling [49] which we adopt in the present work.

III. Evaluation and Results

In this section we turn our attention to the empirical evaluation of the proposed framework. We first succinctly describe our evaluation data set, and then report and discuss the obtained results.

A. Evaluation data

Twitter's Terms of Service explicitly prohibit the sharing or redistribution of tweets, including for research purposes. Consequently, as there was no public benchmark that we could adopt we used a large data set collected by ourselves, acquired as described in [11]. For full detail the reader is referred to the original publication. Here we note the sole difference in the tweet collections used in the present work and that in [11]: their sizes.

Recall that we collected tweets using Twitter's stream and search APIs by periodically retrieving tweets which contain any of the four keywords "autism", "adhd", "asperger", and "aspie" (or any of their derivatives obtained by suffixation). Following the publication of our original work we continued this process and now have at our disposal a data set over an order of magnitude greater than that initially described in [11]. Our corpus now includes 5,650,989 ASD-related tweets collected in the period starting on 26 August 2013 and ending on 1 Oct 2014 (i.e. more than 13 consecutive months).

1) Algorithm parameters: The analytics framework we described in the preceding sections requires the values of a number of free parameters to be set. For the sake of completeness and full reproducibility of the experiments reported in this paper, here we summarize how this was performed and what values were used. Note that this process needs to be performed only once – the framework is fully automatic thereafter. In addition, in the discussion of possible directions for future work in Sec. IV we also outline the ideas we are currently working on which would allow the optimal values of the free parameters to be learnt directly from data. This would make the entire process automatic and remove the need for human input even to the minimal extent described here.

The first of the parameters, introduced in Sec. II-A, is the size of the term dictionary used to represent documents as fixed-size bags of words/terms. It should be noted that in principle the entire set of terms extracted from the document corpus could be used without any negative effects on the quality of the extracted topic information, as the importance of each term is inferred automatically by our HDP-based model. The sole reason for choosing a subset of terms as a dictionary lies in the reduced computational overhead (both time-wise and storage-wise). Following the usual procedure adopted in the literature (for detailed discussion see e.g. [12]) we construct the dictionary using the most frequent terms which account for approximately 90% of the text of the document corpus [18]. In our case this results in 6560 terms. It is insightful to observe the much greater size of this dictionary than the 1500-term one constructed in exactly the same manner in [12] but using main tweet text only. This reflects the more expressive nature of content linked to from tweets compared to the content within tweets themselves.

The remaining free parameters concern the proposed algorithm for temporal tracking of topic structure changes of the document corpus. In particular, to reflect the highly dynamic nature of Twitter we chose short epochs, each spanning three days, with successive epochs exhibiting an overlap of two days. Furthermore, for the pruning of inter-epoch similarity graphs we used the Jaccard similarity threshold of $\tau_{\text{prune}} = 0.5$ which was determined empirically using a small training corpus and qualitative judgement (see Sec. IV for an indication of how this threshold can be learnt directly from data).

B. Results and discussion

We started our analysis of the results obtained with the proposed method by examining the topic structures discovered with different epochs. As in all related work, given that the technical term 'topic' is understood to be a formalization and a proxy for the more abstract notion of a 'topic' in everyday speech, this analysis is inherently subjective in nature as it is not possible to define or indeed obtain an objective ground-truth.

A representative selection of popular topics discovered within an epoch is shown in Fig. 5. This figure shows four representative topics by the proposed method in a single epoch of our ASD tweet dataset. The topics are visualized using the standard world-cloud representation whereby term frequency is proportional to the corresponding font size (note that different colours are merely used for the sake of easier visualization of a large number of terms of different frequencies in limited space; colour encodes no pertinent information about different terms). There are two key observations that are important to make here. The first of these is that the extracted topics are readily interpreted as being meaningful in the context of what is known about autism. For example, Topic 1 can be seen to relate to the persisting myth of a link between childhood vaccination and autism development [23]. Related issues which concern the effects of thimerosal (a vaccine preservative) and mercury which feature in this topic are also related to this myth [30]. Topic 2 concerns Chris Tuttle, a man with Asperger's syndrome working in Wegmans who attracted much media attention in November 2013 after being yelled at by a customer. Topic 3 captures discussions of research on the relationship between autism and various health conditions and reflects the increased willingness of the general public, particularly those affected by rare and debilitating conditions, to examine the scientific literature. Lastly Topic 4 pertains to the schooling of children with learning disabilities which is one of the most common concerns of many parents with autistic children (Lou Salza, whose name features significantly in this topic, is the outspoken headmaster of a school for children with learning disabilities). The second important observation, one which strongly supports the main premise of the present work and demonstrates the usefulness of the proposed framework, is that the topics feature more than a couple of dominant (by frequency) terms. This means that they cannot be readily captured using a small set of keywords such as hashtags or words extracted directly from tweets.

Having demonstrated that our method extracts meaningful topics, we next sought to investigate what type of information was obtained by tracking changes in the topic structure over epochs. Here too we readily observed that the proposed

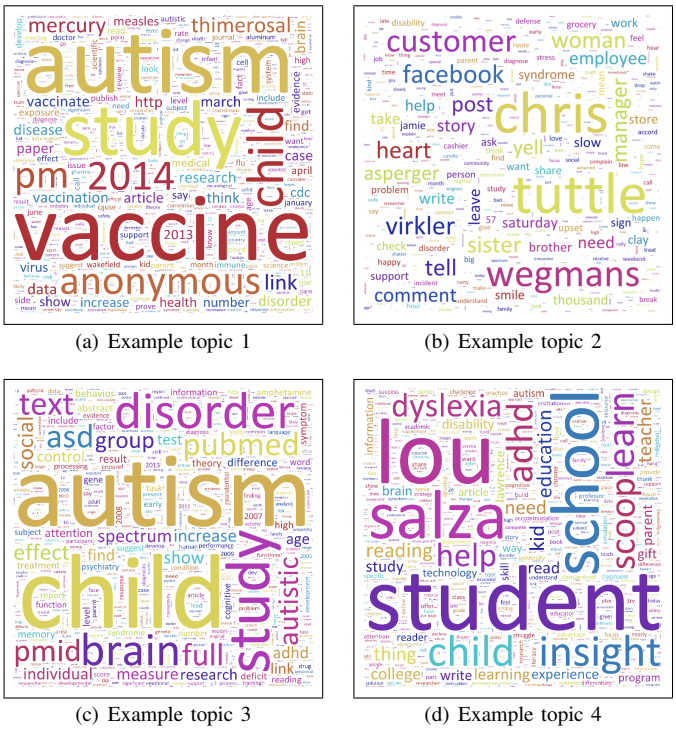

(a) Example topic 1 (b) Example topic 2

(c) Example topic 3 (d) Example topic 4

Fig. 5. An illustration of typical topics discovered by the proposed method in a single epoch of our ASD tweet dataset, shown as size-coded word-clouds – a larger font indicates a proportionally more probable term (different colours are merely used for easier visualization and encode no information pertaining to the corresponding terms themselves). Topic 1 can be readily related to the persistent myth of a link between childhood vaccination and autism development as well as related issues concerning the effects of thimerosal and mercury [30], Topic 2 to the incident involving Chris Tuttle, a Wegmans employee with Asperger's syndrome who attracted media attention in November 2013, Topic 3 to medical literature on autism, children, and brain development, and Topic 4 to the schooling of children with learning disabilities.

method was powerful in extracting interesting and insightful knowledge. Illustrative examples are pictured in Fig. 6. The diagrams in Fig. 6(a) and 6(b) show examples of "linear" topic evolutions i.e. changes in the nature of conversation pertaining to a single, continuously developing topic. For example, the evolution in Fig. 6(a) concerns the bestselling book "The Reason I Jump: The Inner Voice of a Thirteen-Year-Old Boy with Autism" written by the autistic Japanese author Naoki Higashida, and which amongst other things discusses the highly emotive issues of language and intelligence [26]. On the other hand, Fig. 6(c) shows multiple topics merging to form new topics. For example, the topic on behavioural and social aspects of schooling can be seen to merge with the topic concerning the already mentioned myth about a link between vaccination and autism, to produce a topic which encompasses a combination of these issues. Lastly, the ability of our framework to capture and track highly dynamic changes in the topic structure, which is crucial in the analysis of a communication medium such as Twitter, is corroborated quantitatively by the plot in Fig. 7. The plot shows changes in the number of topics per epoch, as well as the corresponding rates of new topic emergence ("birth"), topic disappearance ("death"), merging, and splitting. It can be seen that our framework is effective in capturing the highly dynamic nature of Twitter information exchange as witnessed by the changes in the number of inferred topics per epoch across time, consistently observed emergence of new topics and disappearance of others, topic splitting and merging. It is interesting to observe the high rate of new topic emergence, compared to the rate of topic creation through the merging of old topics, which illustrates the ephemeral nature of most Twitter exchanges [28].

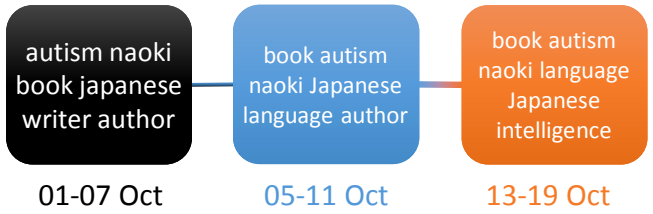

(a) Example topic progression 1 (linear)

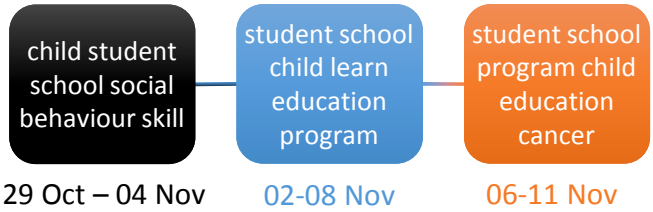

(b) Example topic progression 2 (linear)

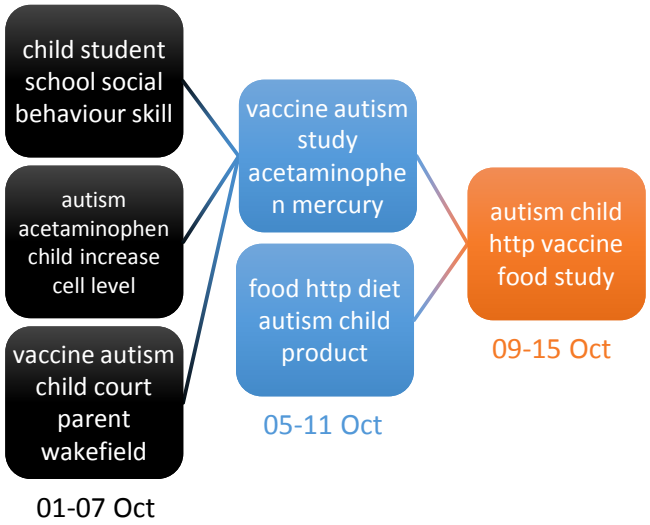

(c) Example topic progression 3 (merging)

Fig. 6. Examples of automatically discovered Twitter topic temporal changes: (a,b) linear evolution of topic content, and (c) topic creation through the merging of topics from the previous epoch. Also see Fig. 7 for related quantitative corroboration of our findings.

IV. Summary and conclusions

In this paper we described the first work on the extraction of the topic structure of tweets and the tracking of its evolution over time. Although our specific aim was that of analysing autism-related Twitter content, the developed framework is entirely general in nature and can be applied to any longitudinal corpus of tweets. Our work identified and elucidated the inherent difficulty associated with topic modelling on tweets – their insufficient information content. To overcome this limitation we proposed an ingenious framework which uses tweets not as endpoint data mining sources but rather as intermediaries used to discover much richer associated content.

The method proposed in this paper opens a rich cornucopia of possible avenues for further research which we intend to pursue. Some of these concern developments of the overall framework. For example in cases when URL targets include videos (say, on YouTube) the use of computer vision or audio processing can add valuable information for even more sophisticated data mining. On the other hand, there are technical aspects of the proposed method which can be improved further. For example we intend to explore the use of ideas from information-theory for fully automatic similarity graph pruning, which would eliminate the need for the free parameter, in the form of the pruning threshold, described in this work. We also intend to explore automatic ways of labelling topics semantically using meaningful sentence fragments by back-analysing probabilistically the collected text data for persistent ngrams across documents with shared topics [5].

References

[1] Autism spectrum disorder fact sheet. *American Psychiatric Publishing*, 2013. 1

[2] F. Abel, Q. Gao, G. J. Houben, and K. Tao. Analyzing user modeling on twitter for personalized news recommendations. *UMAP*, 2011. 3

[3] A. Agarwal, B. Xie, I. Vovsha, O. Rambow, and R. Passonneau. Sentiment analysis of Twitter data. *LSM*, 2011. 3

[4] O. Arandjelović. Computationally efficient application of the generic shape-illumination invariant to face recognition from video. *PR*, 2012. 3

[5] O. Arandjelović. Reading ancient coins: automatically identifying denarii using obverse legend seeded retrieval. *ECCV*, 2012. 8

[6] O. Arandjelović and R. Cipolla. Achieving robust face recognition from video by combining a weak photometric model and a learnt generic face invariant. *PR*, 2013. 3

[7] S. Asur and B. A. Huberman. Predicting the future with social media. *WI-IAT*, 2010. 1

[8] E. Baucom, A. Sanjari, X. Liu, and M. Chen. Mirroring the real world in social media: Twitter, geolocation, and sentiment analysis. *MNLP*, 2013. 1

[9] A. J. Baxter, T. S. Brugha, H. E. Erskine, R. W. Scheurer, T. Vos, and J. G. Scott. The epidemiology and global burden of autism spectrum disorders. *Psychological Medicine*, 2015. 2

[10] A. Beykikhoshk, O. Arandjelović, D. Phung, and S. Venkatesh. Hierarchical Dirichlet process for tracking complex topical structure evolution and its application to autism research literature. *PAKDD*, 2015. 4

[11] A. Beykikhoshk, O. Arandjelović, D. Phung, S. Venkatesh, and T. Caelli. Data-mining Twitter and the autism spectrum disorder: a pilot study. *ASONAM*, 2014. 3, 4, 6

[12] A. Beykikhoshk, O. Arandjelović, D. Phung, S. Venkatesh, and T. Caelli. Using Twitter to learn about the autism community. *SNAM*, 2015. 1, 6

[13] A. Bifet and E. Frank. Sentiment knowledge discovery in Twitter streaming data. *DS*, 2010. 3

[14] D. Blei and J. Lafferty. Correlated topic models. *NIPS*, 2006. 3

[15] J. Bollen, H. Mao, and A. Pepe. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. *ICWSM*, 2011. 3

[16] J. Chang, S. Gerrish, C. Wang, J. L. Boyd-graber, and D. M. Blei. Reading tea leaves: how humans interpret topic models. *NIPS*, 2009. 3

[17] C. Chew and G. Eysenbach. Pandemics in the age of Twitter: content analysis of Tweets during the 2009 H1N1 outbreak. *PLOS ONE*, 2010. 3

[18] A. Clauset, C. R. Shalizi, and M. E. Newman. Power-law distributions in empirical data. *SIAM Review*, 2009. 6

[19] A. Culotta. Towards detecting influenza epidemics by analyzing Twitter messages. *SOMA*, 2010. 1

[20] J. T. Danial and J. J. Wood. Cognitive behavioral therapy for children with autism: Review and considerations for future research. *Journal of Developmental & Behavioral Pediatrics*, 2013. 2

[21] T. S. Ferguson. A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 1973. 5

[22] D. E. Gray. Perceptions of stigma: The parents of autistic children. *Sociology of Health & Illness*, 1993. 2

[23] L. Gross. A broken trust: lessons from the vaccineautism wars. *PLoS Biology*, 2009. 6

[24] J. W. Harrington, L. Rosen, A. Garnecho, and P. A. Patrick. Parental perceptions and use of complementary and alternative medicine practices for children with autistic spectrum disorders in private practice. *Journal of Developmental & Behavioral Pediatrics*, 2006. 2

[25] A. Harshavardhan, A. Gandhe, R. Lazarus, S. H. Yu, and B. Liu. Predicting flu trends using Twitter data. *INFOCOM*, 2011. 1

[26] N. Higashida. *The reason I jump: the inner voice of a thirteen-year-old boy with autism.* Random House, 2013. 7

[27] I. Himelboim and J. Y. Han. Cancer talk on Twitter: community structure and information sources in breast and prostate cancer social networks. *Journal of Health Communication*, 2014. 1

[28] J. Huang, K. M. Thornton, and E. N. Efthimiadis. Conversational tagging in Twitter. *HT*, 2010. 7, 9

[29] C. Hutchings. Commercial use of facebook and twitterrisks and rewards. *Computer Fraud Security*, 2012. 1

[30] A. Hviid, M. Stellfeld, J. Wohlfahrt, and M. Melbye. Association between thimerosal-containing vaccine and autism. *The Journal of the American Medical Association*, 2003. 6, 7

[31] H. Ishwaran and L. F. James. Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 2001. 5

[32] J. W. Jacobson, J. A. Mulick, and G. Green. Costbenefit estimates for early intensive behavioral intervention for young children with autismgeneral model and single state case. *Behavioral Interventions*, 1998. 1

[33] J. Jashinsky, S. H. Burton, C. L. Hanson, J. West, C. Giraud-Carrier, M. D. Barnes, and A. T. Tracking suicide risk factors through Twitter in the US. *Crisis: The Journal of Crisis Intervention and Suicide Prevention*, 2014. 3

[34] L. Jiang, M. Yu, M. Zhou, X. Liu, and T. Zhao. Target-dependent Twitter sentiment classification. *ACL*, 2011. 3

[35] S. E. Levy, D. S. Mandell, and R. T. Schultz. Autism. *Lancet*, 2009. 1

[36] J. Li and C. Cardie. Early stage influenza detection from Twitter. *arXiv preprint*, 2013. 1

[37] J. Lin and D. Ryaboy. Scaling big data mining infrastructure: the Twitter experience. *ACM SIGKDD Explorations Newsletter*, 2013. 1

[38] J. H. Miles. Autism spectrum disorders–a genetics review. *Genetics in Medicine*, 2011. 1

[39] L. Mitchell, M. R. Frank, K. D. Harris, P. S. Dodds, and C. M. Danforth. The geography of happiness: connecting Twitter sentiment and expression, demographics, and objective characteristics of place. *PLOS ONE*, 2013. 3

[40] A. T. Newton, A. D. I. Kramer, , and D. N. McIntosh. Autism online: a comparison of word usage in bloggers with and without autism spectrum disorders. *SIGCHI*, 2009. 3

[41] M. J. Paul and M. Dredze. You are what you tweet: analyzing Twitter for public health. *ICWSM*, 2011. 3

[42] M. J. Paul and M. Dredze. A model for mining public health topics from Twitter. *Health*, 2012. 1

[43] B. Robinson, R. Power, and M. Cameron. An evidence based earthquake detector using Twitter. *LPCI*, 2013. 3

[44] M. A. Russell. *Mining the Social Web: Data Mining Facebook, Twitter, LinkedIn, Google+, GitHub, and More.* O'Reilly Media, Inc., 2013. 1

[45] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes Twitter users: real-time event detection by social sensors. *WWW*, 2010. 3

[46] J. Sankaranarayanan, H. Samet, B. E. Teitler, M. D. Lieberman, and J. Sperling. Twitterstand: news in tweets. *SIGSPATIAL GIS*, 2009. 3

[47] D. Scanfeld, V. Scanfeld, and E. L. Larson. Dissemination of health information through social networks: Twitter and antibiotics. *American Journal of Infection Control*, 2010. 3

[48] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Sharing clusters among related groups: hierarchical Dirichlet processes. *NIPS*, 2004. 5

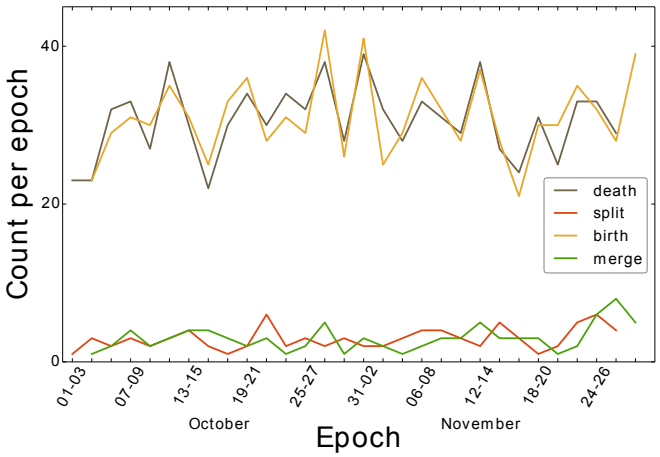

Fig. 7. Complex topic evolution statistics (births, deaths, merges, and splits) over a two month period (Oct–Nov 2015). It can be seen that our framework is effective in capturing the highly dynamic nature of Twitter information exchange as witnessed by the changes in the number of inferred topics per epoch across time, consistently observed emergence of new topics and disappearance of others, topic splitting and merging (also see Fig. 6 for specific examples). Observe the high rate of new topic emergence, compared to the rate of topic creation through the merging of old topics, which illustrates the ephemeral nature of most Twitter exchanges [28].

[49] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 2006. 5

[50] D. Trembath, S. Balandin, and C. Rossi. Crosscultural practice and autism. *Journal of Intellectual and Developmental Disability*, 2005. 2

[51] Twitter. About. *https://about.twitter.com/company*. (accessed April 2015). 1

[52] S. Verma, S. Vieweg, W. J. Corvey, L. Palen, J. H. Martin, M. Palmer, A. Schram, and K. M. Anderson. Natural language processing to the rescue? Extracting "situational awareness" tweets during mass emergency. *ICWSM*, 2011. 3

[53] Z. Warren, M. L. McPheeters, N. Sathe, J. H. Foss-Feig, A. Glasser, and J. Veenstra-VanderWeele. A systematic review of early intensive intervention for autism spectrum disorders. *Pediatrics*, 2011. 2